

AMENDMENTS TO THE SPECIFICATION

Please amend paragraph 15, as follows:

According to an embodiment of the invention, when deciding whether to migrate to a distributed architecture, one could factor in possible increases in overhead costs to discover the true speed-up or scale-up that could be obtained. Workloads which result in greater frequency of data conflicts may not perform well in a distributed architecture because of increased communications costs for maintaining cache consistency. Consequently it is desirable to provide ~~to~~ a mechanism to estimate the number of data conflicts that would occur during execution of a workload in a distributed environment. Additionally, it is useful to determine how to distribute the workload across a cluster so as to reduce the number of data conflicts.

Please amend paragraph 17, as follows:

Fig. 2 depicts the system of Fig. 1 after migration to a multi-node system. For illustrative purposes, this system comprises two separate nodes 202a and 202b, with each node maintaining memories 209a and 209b having caches 210a and 210b, respectively. Note that the clients ~~102~~ 120a-d accessing database 108 through a single node 102 in Fig. 1 has been subdivided among nodes 202a and 202b in Fig. 2. One of the advantages of a distributed, multi-node architecture is that a workload can be distributed across multiple computer nodes, thereby reducing performance bottlenecks and increasing parallelism, load-balancing, and scalability. One way to accomplish these performance advantages is to distribute client accesses across different nodes in the multi-node system. Thus, clients 120a and 120b are shown accessing the first node 202a while clients 120c and 120d are shown accessing the second node 202b. In addition, the work handled through the original

sessions 106 in the single node 102 are also redistributed to sessions 206a and 206b when migrating to the multi-node environment. Thus, sessions S_1 - S_4 shown in node 102 of Fig. 1 have been subdivided such that the work performed by sessions S_1 and S_3 are distributed to node 202a and the work performed by sessions S_2 and S_4 are distributed to node 202b.

Please amend paragraph 27, as follows:

To illustrate, the above trace record can also be used to calculate the number of data conflicts if the workload is migrated to a 3-node system. Assume that the sessions on single node 102 are subdivided for a 3-node system using modulo division such that a first modulo class includes sessions S_1 and S_4 , a second modulo class includes sessions S_2 , and a third modulo class includes session S_3 . As noted above, it is assumed that the sessions likely to be running in the same node in a multi-node system also belong to the same modulo class. Thus, sessions S_1 and S_4 correspond to a first node, session S_2 corresponds to a second node, and session S_4 S_3 corresponds to a third node.